

Using Neural Network Formalism to Solve Multiple-Instance Problems

Tomáš Pevný^{1,2} and Petr Somol¹

¹Cisco Systems, Charles Square 10, Prague 2, Czech Republic

²Czech Technical University, Faculty of Electrical Engineering,
Charles Square 13, Prague 2, Czech Republic

Abstract. Many objects in the real world are difficult to describe by a single numerical vector of a fixed length, whereas describing them by a set of vectors is more natural. Therefore, *Multiple instance learning* (MIL) techniques have been constantly gaining on importance throughout last years. MIL formalism represents each object (sample) by a set (bag) of feature vectors (instances) of fixed length where knowledge about objects (e.g., class label) is available on bag level but not necessarily on instance level. Many standard tools including supervised classifiers have been already adapted to MIL setting since the problem got formalized in late nineties. In this work we propose a neural network (NN) based formalism that intuitively bridges the gap between MIL problem definition and the vast existing knowledge-base of standard models and classifiers. We show that the proposed NN formalism is effectively optimizable by a modified back-propagation algorithm and can reveal unknown patterns inside bags. Comparison to eight types of classifiers from the prior art on a set of 14 publicly available benchmark datasets confirms the advantages and accuracy of the proposed solution.

1 Motivation

The constant growth of data sizes and data complexity in real world problems has increasingly put strain to traditional modeling and classification techniques. Many assumptions cease to hold; it can no longer be expected that a complete set of training data is available for training at once, models fail to reflect information in complex data unless a prohibitively high number of parameters is employed, availability of class labels for all samples can not be realistically expected, and particularly the common assumption about each sample to be represented by a fixed-size vector seems to no longer hold in many real world problems.

Multiple instance learning (MIL) techniques address these concerns by allowing samples to be represented by an arbitrarily large set of fixed-sized vectors instead of a single fixed-size vector. Any explicit ground truth information (e.g., class label) is assumed to be available on the (higher) level of samples but not on the (lower) level of instances. The aim is to utilize unknown patterns on instance-level to enable sample-level modeling and decision making. Note that MIL does not address the Representation Learning problem [3]. Instead it aims at better

utilization of information in cases when ground truth knowledge about a dataset may be granular and available on some of the various levels of abstraction only.

From practical point of view MIL promises to i) save ground truth acquisition cost - labels are needed on sample-level, i.e., on higher-level(s) of abstraction only, ii) reveal patterns on instance-level based on the available sample-level ground truth information, and eventually iii) achieve high accuracy of models through better use of information in data.

Despite significant progress in recent years, the current battery of MIL tools is still burdened by compromises. The existing models (see next Section 2 for discussion) clearly leave open space for more efficient utilization of information in samples and for clearer formalism to provide easily interpretable models with higher accuracy. We aim at providing a formalism that effectively bridges the gap between MIL problem formulation and standard battery of modeling and classification techniques.

2 Multi-instance problem

The pioneering work [9] coined *multiple-instance* or *multi-instance* learning as a problem, where each sample b (to be denoted *bag* in the following) consists of a set of instances x , i.e., $b = \{x_i \in \mathcal{X} | i \in \{1, \dots, |b|\}\}$, $b \in \mathcal{B} = \cup_{k \geq 1} \cup \{x_i \in \mathcal{X} | i \in \{1, \dots, k\}\}$, where each instance x can be attributed a label $y_x \in \{-1, +1\}$, but these instance-level labels are not known even in the training set. The sample b was deemed positive, if at least one of its instances had a positive label, i.e., label of a sample b is $y = \max_{x \in b} y_x$. For this scenario the prevalent approach is the so-called *instance-space paradigm*, i.e., to train a classifier on the level of individual instances $f : \mathcal{X} \mapsto \{-1, +1\}$ and then infer the label of the bag b as $\max_{x \in b} f(x)$. See Sect. 2.1 for overview of instance-space paradigm methods. Neural networks have been previously applied to MIL problems in instance space paradigm (see Sect. 2.4).

Later works (see reviews [1,11]) have introduced different assumptions on relationships between the labels on the instance level and labels of bags or even dropped the notion of instance-level labels and considered only labels on the level of bags, i.e., it is assumed that each bag b has a corresponding label $y \in \mathcal{Y}$, which for simplicity we will assume to be binary, i.e., $\mathcal{Y} = \{-1, +1\}$ in the following.

The common approach of the latter type is either to follow a *bag-space paradigm* and define a measure of distance (or kernel) between bags (see Sect. 2.2 for overview of existing methods) or to follow an *embedded-space paradigm* and define a transformation of the bag to a fixed-size vector (see Sect. 2.3 for overview).

This paper adopts a definition inspired by Support-Measure Machines [18] assuming that probability distributions of instances in positive and negative bags are different and the probability distributions generating instances in each bag are different. Based on this assumption we introduce in Sect. 3 the key idea of this paper - a Neural Network formalism that effectively embeds bags into a vector space, providing a generic solution to the MIL problem.

2.1 Instance-space paradigm

Prior art algorithms falling into this category assume that every positive bag contains at least one (positive) instance, which does not occur in instances of negative bags. Therefore these algorithms assume a labeling on instances, even though such labeling is not known.

EM-DD[23] fits a generative model differentiating between positive and negative instances using an EM algorithm to infer the unknown labels of instances in positive bags in the E-step. In M-step it maximizes the diverse density [17], which measure for point x how many different positive bags have an instance near x and how far the negative instances are from x . Intuitively, the diversity density of a hypothesis h is just the likelihood (with respect to the data) that h is the target. A high diverse density indicates a good candidate for a "true" concept.

MILBoost[22] adapts boosting to the MIL problems. The proposed adaptation first converts the real-valued outputs of the weak classifier of choice to probabilistic values using sigmoid function $p_{ij} = \frac{1}{1+\exp(-z_{ij})}$. Then, to predict the class of a bag the instance-level verdicts are combined by means of soft-or defined as $p_i = 1 - \prod_{j \in b_i} (1 - p_{ij})$. The algorithm finally uses boosting to optimize the likelihood of training bags $L = \prod_i p_i^{y_i} (1 - p_i)^{1 - y_i}$, which is now possible since the function is differentiable.

MI-SVM [2] modifies the standard formulation of SVM to include the instance-space assumption that each positive bag should contain at least one positive instance, as

$$\arg \min \frac{\lambda}{2} \|w\|^2 + \sum_{i=1}^l \xi_i$$

subject to

$$\begin{aligned} - < w, x > &\geq 1 - \xi_i, \forall x \in b_i \wedge y_i = -1 \\ < w, s(b) > &\geq 1 - \xi_i y_i = +1, \end{aligned}$$

where $s(b) = \arg \max_{x \in b} < w, x >$ denotes the pattern selected as the positive "witness" in b_i and $f(b) = \max_{x \in b} < w, x >$ is the output of the classifier.¹

2.2 Bag-space paradigm

Methods in bag-space paradigm either define a kernel function $k : \mathcal{B} \times \mathcal{B} \mapsto \mathbb{R}_0^+$ and use a kerneled classifier or they define a distance function and take use of the nearest-neighbor principle.

Citation-kNN [21] uses Hausdorff distance to calculate distance between two bags. The classification of b is based on the difference between number of positive "referencing" and "citing" nearest neighbors and negative ones. Referencing nearest-neighbor bags to b are the usual k -nearest neighbors, the citing nearest neighbors are those to which b is the nearest neighbor.

¹ To simplify the notation we have dropped the bias from the formulation.

Probably the first kernel defined on sets was the convolutional kernel introduced in [14], but it was used on MIL problems. Ref, [12] uses this kernel as a departure point, but then proposes to embed the bag into the fixed-dimensional Euclidean space, and therefore we consider the method to belong to bag-space paradigm. A recent work [18] utilizes universal kernels defined over sets [20] to solve general MIL problems (as defined later in Sect. 3) by kernelized Support Vector Machines. The biggest advantage of this work is asymptotical optimality, though the complexity prevents it to be used on large problems.

2.3 Embedded-space paradigm

Methods from embedded-space paradigm represent the whole bag as a fixed-length feature vector, thus enabling subsequent use in any off-the shelf algorithm. Formally, bag b is embedded into a m -dimensional vector as

$$(\phi_1(b), \phi_2(b), \dots, \phi_m(b)) \in \mathbb{R}^m, \quad (1)$$

where ϕ is an embedding function $\phi : \mathcal{B} \mapsto \mathbb{R}$.

The basic methods [12,21] calculate a statistics such as mean, maximum, minimum, or covariance of each individual feature of instance vectors over the bags. Thus each bag is represented by a feature vector of the dimension equal to that of instance feature vectors. It is both an advantage and a drawback that the mapping is not in any sense adapted to the solved problem.

Miles [4] implements embedding $\phi : \mathcal{B} \mapsto \mathbb{R}$ using an instance-level dictionary $\mathcal{D} = \{d_i \in \mathcal{X} | i \in \{1, \dots, m\}\}$ such that

$$\phi = \sum_{x \in b} k(x, d)$$

with k being

$$k(x, d) = \begin{cases} \exp^{-\frac{1}{\sigma^2} \|x-d\|^2} & \text{iff } d \text{ is the nearest neighbor of } x \text{ in } \mathcal{D} \\ 0 & \text{otherwise,} \end{cases}$$

The choice of dictionary items is approached as a feature-selection problem and solved by a series of 1-norm Support Vector Machines, from which stems its high computational complexity.

A general approach to (1) was recently proposed in [7] defining ϕ using a distance function (or kernel) over the bags, i.e. $k : \mathcal{B} \times \mathcal{B} \mapsto \mathbb{R}$, and dictionary \mathcal{D} containing whole bags instead of individual instances. The proposed kernels include Hausdorff distance and various combination of means, maxima, and minima over instances in the kernel. Note that this approach can be thought about as a crude approximation of a kernel defined over the space of bags [18].

2.4 Instance-space neural network

The current adaptations of neural networks to MIL fall into the instance-space paradigm. Ref. [19] proposes a smooth approximation of the maximum pooling in the last neuron as

$$\frac{1}{|\mathcal{B}|} \ln \left(\sum_{x \in b} \exp(f(x)) \right),$$

where $f(x) : \mathcal{X} \mapsto \mathbb{R}$ is the output of the network before the pooling. Ref. [24] drops the requirement on smooth pooling and uses the maximum pooling function in the last neuron. Both approaches optimize the L_2 norm error function.

3 Neural network formalism

The formalism proposed below is intended for a general formulation of MIL problems. It assumes existence of a space of probability distributions $\mathcal{P}^{\mathcal{X}}$ where each distribution $p \in \mathcal{P}^{\mathcal{X}}$ is assumed to generate instances in some bag and is defined over the instance space \mathcal{X} . Each bag b is assumed to be a realization of a probability $P(p_b, y)$, where $y \in \mathcal{Y}$ is the bag label and p_b is the probability distribution on instances $x \in b$. During the learning process each concrete bag b is thus viewed as a realization of unknown probability distribution p_b that can only be inferred from groups of instances $\{x \in b | x \sim p_b\}$ observed in data. The goal is to learn a discrimination function $f : \mathcal{B} \mapsto \mathcal{Y}$, where \mathcal{B} is the set of all possible realizations of distributions $p \in \mathcal{P}^{\mathcal{X}}$, i.e., $\mathcal{B} = \{x_i | p \in \mathcal{P}^{\mathcal{X}}, x_i \sim p, i \in \{1, \dots, l\}, l \in \mathbb{N}\}$. This definition includes that used in [9], which corresponds to the assumption that there exist instances which are never generated by distributions of negative samples, which means that the difference of support of positive and negative probability distributions is non-empty, i.e., $p_+ \setminus p_- \neq \emptyset$, where $p_+ \sim P(p|+)$ and $p_- \sim P(p|-)$. On the other hand it also includes the case, where every instance can occur in positive and negative samples, but some instances are more frequent in the negative samples or in the positive samples.

Regarding the choice of MIL paradigm we argue that there is only little difference between *embedded-space* and *bag-space* paradigms, because (i) many machine learning methods that accept fixed-length feature vectors can be used with kernels (or kernels on top of kernels [18,20]); (ii) kernels can be viewed as an embedding $\phi : \mathcal{B} \mapsto \mathcal{H}$ into some unknown Hilbert space \mathcal{H} with a dot-product such that $k(b_i, b_j) = \langle \phi(b_i), \phi(b_j) \rangle_{\mathcal{H}}$; (iii) and finally kernel functions can be approximated by a dot-product in Euclidean space [7,10], which effectively allows the kernel methods to scale.

In the rest of this section, for simplicity and the reasons above we follow the *embedded-space paradigm*, which represents the bag b in an m -dimensional Euclidean space \mathbb{R}^m through set of mappings

$$(\phi_1(b), \phi_2(b), \dots, \phi_m(b)) \in \mathbb{R}^m. \quad (2)$$




Fig. 1. Neural network optimizing the embedding in embedding-space paradigm.

Many existing methods implement embedding function $\phi : \mathcal{B} \mapsto \mathbb{R}$ as

$$\phi = g \left(\{k(x, d)\}_{x \in b} \right), \quad (3)$$

where $k : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}_0^+$ is a suitably chosen distance function, $g : \cup_{n=1}^{\infty} \mathbb{R}^k \mapsto \mathbb{R}$ is the pooling function (e.g. minimum, mean or maximum), and finally $\mathcal{D} = \{d_i \in \mathcal{X} | i \in \{1, \dots, m\}\}$ is the dictionary with instances as items. Various existing methods differ in the choice of aggregation function g , distance function k , and finally in selection of dictionary items $d \in \mathcal{D}$. A generalization was recently proposed in [7] defining ϕ using a distance function (or kernel) over the bags $k : \mathcal{B} \times \mathcal{B} \mapsto \mathbb{R}$ and dictionary \mathcal{D} containing bags rather instances. This generalization can be seen as a crude approximation of kernels over probability measures used in [18,20].

To optimize projection functions $\{\phi_i\}_{i=1}^m$, which we believe to differ between problems, we think about the computational model in Equation (1) as about a neural network. One (or more) lower layers followed by a pooling layer implement the embedding function ϕ , and subsequent layers implement the classifier that already uses representation of the bag as a feature vector of fixed length. The model is sketched in Figure 1 with a single output neuron implementing a linear classifier once the embedding to a fixed-length feature representation is done. Adopting the neural network formalism enables to optimize individual components of the embedding function as follows.

- Lower layers (denoted in Figure 1 as $\{k_i\}_{i=1}^m$) before pooling identifies parts of the instance-space \mathcal{X} , where the probability distributions generating instances in positive and negative bags differs the most with respect to the chosen pooling operator.
- The pooling function g can be either fixed to minimum or maximum, or it can be optimized for example as in [13], where the pooling function is imple-

- mented as $\sqrt[q]{\frac{1}{|b|} \sum_{i \in b} |x_i|^q}$ with the parameter q optimized for each neuron separately by back-propagation. Although the optimization of q is appealing, our experiments with the method of [13] did not yield any improvement.
- Layers after the pooling (denoted in Figure 1 as $f(\cdot)$) optimizes the classifier that already uses the fixed representation.

The above model is very general and allows automatic optimization of all parameters by means of back-propagation, though the user still needs to select the number of layers, number of neurons in each layer, their transfer function, and possibly pooling function. Concrete choices and their effect on the error are discussed in Experimental section 4.2.

Remark: the key difference of this approach to those described in Sect. 2.4 is in performing pooling in one or multiple layers inside the network as opposed to after the last neuron or layer. This difference is key to the shift from instance-centric modeling in prior art to bag-centric advocated here.

4 Experimental evaluation

Datasets *Biocreative component*, *Birds*, *Corel african*, *Elephant*, *Fox*, *Tiger*, *Mesidor*, *Musk1*, *Musk2*, *Mutagenesis easy*, *Mutagenesis hard*, *Newsgroups*, *Breast*, and *Web* used to evaluate the above formalism were real-life datasets downloaded from <http://mipproblems.org> used in [6,7]. For almost all problems bag-space or embedded-space paradigms were more appropriate, except *Musk* dataset used in [9] and *Newsgroup*, which can be considered as belonging to instance-space paradigms.

4.1 Comparison to prior art

Our first experiment compares the proposed formalism to the prior art. We compare our results to those of the evaluation published in [7] on 22 MIL classifiers. From the referred list we include in our evaluation only those classifiers the lowest equal error rate (EER) on at least one real-life dataset.² The fact that this selection yielded to 12 classifiers for 14 problems demonstrates how diverse the MIL problems are and how difficult it is to choose suitable classifier for the given problem. The compared classifiers include representatives of *instance-space paradigm*: EM-DD [23], MIL Boost [22], SimpleMIL which assumes all instances within the bag to share the bag’s label, and MI-SVM [2] with Gaussian kernel; *bag level paradigm*: k -nearest neighbor with citation distance[21] using 5 nearest neighbors; and finally *embedded space paradigm*: Miles [4] with Gaussian kernel, Bag dissimilarity [7] with *minmin*, *meanmin*, *meanmean*, and *Hausdorff distance*, and Cov-coef [5] calculating covariance coefficient of each feature over the bag. All four variants of Bag dissimilarity embedding, Cov-coef embedding, and

² The equal error rate used in [5] is optimistic, since it was calculated as a minimum of an average of false positive and false negative rate on the testing set, where the minimum was taken over all possible choices of the threshold.




Fig. 2. Critical difference diagram of the compared methods from prior art and the proposed formalism that uses neural network with two layers, each having 20 ReLu neurons followed by max aggregation function and linear output unit.

Simple MIL methods used logistic regression as classifier for the fixed-feature vector representation.

The proposed neural network MIL classifier used for all problems the same topology. It consists of two layers each having 20 neurons with a rectified linear (ReLu) [15] transfer function ($\max\{0, x\}$), followed by the max-pooling layer and finally followed by a single linear output unit. The network was optimized to minimize average hinge loss function $\max\{0, 1 + y(f(b))\}$, where $f(b) : \mathcal{B} \mapsto \mathbb{R}$ is the function implemented by the neural network. The network was optimized by Adam [16] allowed to run for 10 000 iterations with default settings as suggested in the original publication and the gradient being estimated from all available samples. The regularization on network weights was L_1 with 10^{-8} .

The measure used for comparison of classifiers is the equal error rate (EER) calculated as the average of false positive and false negative rates estimated by five times repeated 10-fold stratified cross-validation, which is the same approach as used in [7]. The same measure together with exactly same the same testing sets allowed us to directly refer to results on the prior art from the aforementioned publications. EER together with the best algorithm from the prior art is summarized in Table 1 and EER of all compared classifiers is deferred to Table 2 in Appendix. The results show that using the neural network formalism to optimize the embedding function (3) and the linear classifier appears superior to all prior art. (Note that we did not employ any optimization of the pooling function and the chosen architecture is very sub-optimal as is shown below.) Figure 2 summarizes the results in critical difference diagram [8] showing the average rank of each compared classifier. The confidence interval of corrected Bonferroni-Dunn test with significance 0.05 of whether two classifiers have equal overall performance is shown in thick black line. We observe that the classifier implemented using the proposed neural net formalism achieved the best performance and is the only classifier that is statistically better than the rest.

problem	EER of the neural net	EER of the best algorithm of prior art
Biocreative component	8.0	25.0 MI SVM
Birds	5.0	11.2 MIL Boost
Corel african	5.1	11.2 Bag dis. with minmin
Elephant	5.0	16.3 Bag dis. with minmin
Fox	7.5	36.0 EM-DD
Tiger	6.0	19.0 Miles
Messidor	10.1	29.7 Simple MIL
Musk 1	3.2	12.8 KNN
Musk 2	2.9	11.8 Bag dis. with Hausdorff
Mutagenesis hard	11.9	20.8 Bag dis. with minmin
Mutagenesis easy	16.7	16.9 Cov-coef embedding
Newsgroups	5.0	18.4 Bag Dis. with meanmean
Breast	8.9	13.6 MI SVM
Web	5.9	20.9 Miles

Table 1. Column captioned “EER of the neural net” shows the equal error rate estimated by the ten-fold cross-validation of the proposed formalism that uses neural network with two layers, each having 20 ReLu neurons followed by max aggregation function and linear output unit. Column captioned “EER of the best algorithm of prior art” shows the same quantity achieved by the best algorithm from the prior art.

4.2 Topology of the network

In the previous section all neural networks for all problems had the same topology, which was certainly far from being optimal. The rationale behind was to show that despite this weakness, neural networks can optimize the embedding function (3) in accord with the classifier better than fixed methods used in the prior art. This section explores the effect on EER of different topologies by varying (i) number of layers in the network by evaluating networks with one or two layers having 20 ReLu neurons; (ii) pooling functions by testing maximum and mean aggregation function; and finally (iii) the position of the pooling which is either behind the second ReLu layer or after the last linear neuron. The network with maximum pooling after the last linear layer implements the prior art of [24], although here a different loss function is used.

EER of networks with eight different topologies is shown in Table 2 with summary in the form of critical difference diagram in Figure 3. The biggest surprise is that the maximum pooling function, which is presently preferred in convolutional neural networks was inferior to the mean pooling. But thinking about the MIL problems as about classifying probability density measures, the output of the embedding operation (2) when mean pooling is used can be understood as a generalization of a histogram, the outcome of the comparison makes sense. Contrary, the maximum pooling function can be thought about as a relaxed version of an indicator function checking presence of a specific pattern. The experimental results support this reasoning, since the max pooling function performed the best on the “musk 2” dataset, which was a dataset used to define instance-level




Fig. 3. Critical difference of the comparison of neural network architectures.

	pooling before linear unit				pooling after linear unit			
	two layers		one layer		two layers		one layer	
	max	mean	max	mean	max	mean	max	mean
Biocreative component	8.0	7.6	12.0	15.9	24.3	7.8	20.5	14.0
Birds	5.0	3.9	4.2	3.0	10.0	2.9	7.4	2.0
Corel african	5.1	4.8	8.6	9.9	7.2	3.7	16.5	9.2
Elephant	5.0	5.5	1.5	2.5	7.5	4.5	2.5	3.0
Fox	7.5	5.5	10.0	15.0	14.5	10.5	7.5	7.0
Tiger	6.0	3.0	4.0	3.5	5.5	2.5	4.0	3.5
Messidor	10.1	9.4	5.7	21.4	14.7	9.8	9.5	24.3
Musk 1	3.2	0.0	1.1	6.5	2.1	2.2	5.5	0.0
Musk 2	2.9	4.2	1.6	4.2	10.6	3.7	3.4	2.1
Mutagenesis hard	11.9	11.1	13.5	8.8	44.5	11.5	35.3	9.5
Mutagenesis easy	16.7	18.8	24.4	21.0	34.2	12.9	15.4	20.6
Newsgroups	5.0	4.0	6.0	3.0	31.0	4.0	5.0	6.0
Breast	8.9	0.0	5.8	0.0	5.0	0.0	0.0	3.8
Web	5.9	5.9	7.3	25.0	10.5	2.2	12.7	9.3

Table 2. Equal error rate of eight different topologies of neural networks.

paradigm in [9]. These results also demonstrate that methods optimizing the pooling function are of great need, because one can easily imagine that in complicated domains a combination of mean and maximum pooling function can be of great use.

Also observe that the position of the pooling, which is the differentiating factor to the prior art, is important, since topologies mimicking the prior art [24] with pooling after the output linear neuron are clearly inferior to their counterparts with pooling before the output linear layer (the only exception is the network with two layers and maximum pooling function).

Lastly, on average better error of networks with two layers and mean pooling proved that higher flexibility of embedding function is an advantage, though it might lead to overfitting.

5 Conclusion

This work has presented a straightforward generalization of neural networks to multi-instance problems. Unlike the prior art, the proposed formalism embeds samples consisting of multiple instances into vector space, enabling subsequent use with standard decision-making techniques. The key advantages of the proposed solution is that it simultaneously optimizes the classifier and the embedding. We have illustrated these advantages on a set of real-world examples, comparing our results to a large number of algorithms from the prior art. The proposed formalism seems to notably outperform the majority of standard MIL methods in terms of accuracy. The overwhelming success in experimental evaluation shows the expressive power of the proposed formalism as well as its efficiency in training. It should be stressed though that we compare to results published by the respective method authors or by authors of survey benchmarks; not all methods in referred tests may have been set in the best possible way. However, as many such cases would be due to too excessive computational complexity of the respective methods, the formalism we propose becomes competitive also due to its relatively modest computational complexity that does not exceed that of a standard 3-layer neural network. The proposed formalism opens up a variety of options for further development. Better and possibly more automated choice of pooling functions is one of the promising ways to improve performance on some types of data.

References

1. Jaume Amores. Multiple instance classification: Review, taxonomy and comparative study. *Artificial Intelligence*, 201:81–105, 2013.
2. Stuart Andrews, Ioannis Tsochantaridis, and Thomas Hofmann. Support vector machines for multiple-instance learning. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 577–584. MIT Press, 2003.
3. Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *arXiv preprint arXiv:1206.5538v2*, 2012.
4. Yixin Chen, Jinbo Bi, and J. Z. Wang. Miles: Multiple-instance learning via embedded instance selection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(12):1931–1947, Dec 2006.
5. Veronika Cheplygina and David M. J. Tax. Supplemental document to characterizing multiple instance datasets.
6. Veronika Cheplygina and David M. J. Tax. *Similarity-Based Pattern Recognition: Third International Workshop, SIMBAD 2015, Copenhagen, Denmark, October 12-14, 2015. Proceedings*, chapter Characterizing Multiple Instance Datasets, pages 15–27. Springer International Publishing, Cham, 2015.
7. Veronika Cheplygina, David M.J. Tax, and Marco Loog. Multiple instance learning with bag dissimilarities. *Pattern Recognition*, 48(1):264 – 275, 2015.
8. Janez Demšar. Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine Learning Research*, 7:1–30, 2006.

9. Thomas G Dietterich, Richard H Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1):31–71, 1997.
10. Petros Drineas and Michael W Mahoney. On the nyström method for approximating a gram matrix for improved kernel-based learning. *The Journal of Machine Learning Research*, 6:2153–2175, 2005.
11. James Foulds and Eibe Frank. A review of multi-instance learning assumptions. *The Knowledge Engineering Review*, 25(01):1–25, 2010.
12. Thomas Gärtner, Peter A Flach, Adam Kowalczyk, and Alexander J Smola. Multi-instance kernels. In *ICML*, volume 2, pages 179–186, 2002.
13. Caglar Gulcehre, Kyunghyun Cho, Razvan Pascanu, and Yoshua Bengio. Learned-norm pooling for deep feedforward and recurrent neural networks. In *Machine Learning and Knowledge Discovery in Databases*, pages 530–546. Springer, 2014.
14. David Haussler. Convolution kernels on discrete structures. Technical report, University of California at Santa Cruz, 1999.
15. K. Jarrett, K. Kavukcuoglu, M. Ranzato, and Y. LeCun. What is the best multi-stage architecture for object recognition? In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 2146–2153, Sept 2009.
16. Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
17. Oded Maron and Tomás Lozano-Pérez. A framework for multiple-instance learning. In M. I. Jordan, M. J. Kearns, and S. A. Solla, editors, *Advances in Neural Information Processing Systems 10*, pages 570–576. MIT Press, 1998.
18. Krikamol Muandet, Kenji Fukumizu, Francesco Dinuzzo, and Bernhard Schölkopf. Learning from distributions via support measure machines. In *Advances in neural information processing systems*, pages 10–18, 2012.
19. Jan Ramon and Luc De Raedt. Multi instance neural networks. 2000.
20. Bharath K. Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Bernhard Schölkopf, and Gert R.G. Lanckriet. Hilbert space embeddings and metrics on probability measures. *J. Mach. Learn. Res.*, 11:1517–1561, August 2010.
21. Jun Wang and Jean-Daniel Zucker. Solving the multiple-instance problem: A lazy learning approach. In *Proceedings of the Seventeenth International Conference on Machine Learning*, ICML '00, pages 1119–1126, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc.
22. Cha Zhang, John C. Platt, and Paul A. Viola. Multiple instance boosting for object detection. In Y. Weiss, B. Schölkopf, and J. C. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 1417–1424. MIT Press, 2006.
23. Qi Zhang and Sally A. Goldman. Em-dd: An improved multiple-instance learning technique. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, pages 1073–1080. MIT Press, 2002.
24. Zhi-hua Zhou and Min-ling Zhang. Neural networks for multi-instance learning. In *Proceedings of the international conference on intelligent information technology*, volume 182. Citeseer, 2002.

id		NN										Bag dissimilarity		
		EM-DD	MIL	Boost	KNN	SimpleMil	Cov-coef	MISVM	Miles	minmin	meanmin	mean	hausd	
329	Biocreative component	3.7	30.0	27.4	31.2	25.3	42.2	25.0	39.3	63.1	32.5	—	—	
155	Birds	4.9	13.4	11.2	14.4	20.4	13.4	15.0	11.2	17.8	11.9	16.7	12.8	
110	Corel african	0.7	16.1	18.4	17.9	29.6	21.5	18.7	43.7	11.2	11.6	16.2	32.4	
303	Elephant	5.0	21.0	17.4	19.3	21.8	18.2	16.8	18.6	16.3	19.0	19.9	34.5	
301	Fox	7.5	36.0	41.5	40.4	46.6	39.7	36.2	36.0	39.0	36.1	42.4	46.4	
302	Tiger	6.0	31.0	23.0	24.6	26.4	29.6	20.4	19.0	26.6	19.0	29.0	37.0	
333	Messidor	10.0	—	32.7	37.7	29.7	30.8	35.6	35.0	35.6	—	—	—	
101	Musk 1	3.3	18.2	31.4	12.8	23.1	27.1	13.6	14.4	15.2	14.5	21.8	16.6	
102	Musk 2	2.9	19.9	32.4	18.7	28.9	25.0	15.3	12.7	17.0	13.8	20.9	11.8	
315	Mutagenesis hard	10.6	41.5	47.0	44.7	48.0	37.8	41.3	51.7	20.8	35.8	45.0	57.7	
314	Mutagenesis easy	14.3	41.2	23.6	20.4	47.0	16.9	44.9	29.4	26.4	20.0	33.6	22.2	
164	Newsgroups	5.0	52	—	—	23.2	29.6	54.4	51.2	48.0	19.8	18.4	23.4	
332	Breast	8.6	—	23.0	34.4	47.9	14.4	13.6	29.4	23.7	22.3	—	—	
155	Web	2.7	—	28.8	48.5	29.6	25.7	21.3	20.9	41.8	50.9	48.5	—	

Table 3. Equal error rates of all compared classifiers of prior art (copied from [5]) and neural network (NN) used in the comparison in Sect. (4.1) with two layers of 20 ReLu neurons and maximum pooling before the last linear output neuron. The column captioned “id” is the ID of the problem as used in [5]. Fields with missing values are missing due to excessive computational requirements.